Rule WLM007: MSO Service Coefficient may be too large

Finding: CPExpert has determined that the MSO service coefficient in the service definition may be too large.

Impact: This finding should be viewed as generally having a LOW IMPACT on the performance of your computer system. However, under certain circumstances (explained below), this finding could have a HIGH IMPACT on the performance of specific application systems, and may have a HIGH IMPACT on overall system performance.

Logic flow: This a basic finding. There are no predecessor rules.

Discussion: The MSO service coefficient specification in the service definition indicates the value by which accumulated storage service units will be multiplied (or weighted) before the storage service units are included in the accumulated service associated with an address space.

The amount of service consumed by an address space is computed by the formula:

service = (**CPU** x CPU Service Units)
+ (**IOC** x I/O Service Units)
+ (**MSO** x Storage Service Units)
+ (**SRB** X SRB Service Units)

The CPU, IOC, MSO, and SRB values highlighted in the above formula relate to the corresponding values specified for the CPU, IOC, MSO, and SRB service coefficients specified in the Workload Manager ISPF panel. The default values for these coefficients are CPU=10, IOC=5, MSO=0, and SRB=10.

The basic intent of the service consumption algorithm is to create a single value which corresponds to the rate at which an address space consumes system resources. As this rate increases, an installation can chose to assign the address space a lower priority with respect to other address spaces. One result of this approach is that the service consumed by address spaces can be controlled: important work can be given access to system resources and the access of less important work can be restricted. This control is accomplished between domains and within domains.

With MVS/ESA SP5.1, the primary use of the accumulated service associated with an address space is to determine whether the address space is moved to a different service class period.

When the address space accumulates more than a specified amount of service, the SRM will reassign the address space to the next lower performance period (if multiple periods are defined)[1]. The amount of service controlling when an address space is reassigned to the next lower performance period is specified by the value of the **DUR** keyword in the definition of the service class performance period.

The normal purpose of defining multiple performance periods is to give higher priority (smaller response goals or higher importance) to interactive transactions, short batch job steps, etc. Overall response is decreased (and overall throughput is increased) when address spaces requiring relatively few resources are processed at a higher priority than those address spaces requiring substantial resources.

Of particular importance to this finding is the way in which Storage Service Units are computed. The Storage Service Units are computed as:

$$\textit{Storage Service Units} = \frac{(\textit{real page frames})\ x\ (\textit{CPU service units})}{50}$$

The intent of dividing by 50 was to cause the Storage Service Units value to be scaled comparable to the other service units. This intent worked fairly well for systems with a relatively small amount of central storage.

However, the Storage Service Unit values can become quite large and can be unpredictable in systems with a significant amount of central storage. These two potential effects are discussed separately:

- The Storage Service Unit values can become large in systems with a significant amount of central storage. An address space can be allocated a large amount of central storage, so that the Storage Service Unit values completely dominate the Service Unit computation.

[1]Note that the address space is not reassigned immediately upon accumulation of the specified amount of service. The reassignment is performed only when the SRM evaluates the address space for changed conditions. Therefore, the SRM cannot determine when the **exact** amount of service specified in the DUR parameter is used; the SRM can only determine when the DUR value is exceeded.

The SRM checks for transactions exceeding their DUR value only when the SRM gains control. Each address space normally is evaluated every three SRM seconds, and reassignment is performed only if the address space has accumulated more than the value specified for the DUR keyword. Consequently, an address space will always accumulate **more** service than specified by the DUR keyword before being reassigned.

For example, an address space may require an allocation of 500 page frames of central storage. Suppose that the address space used 10 milliseconds of CPU (total TCB and SRB) time per second and issued 10 I/O operations per second. If the default service coefficients were used and the processor was a 3090(Model 600J single image), the service rate would be calculated as:

$$service = (10 * 0.01 * 941.53) + (5 * 10) + \frac{(3 * 500 * 0.01 * 941.53)}{50}$$

$$service = 94.153 + 50 + 282.459$$

$$service = 426.612$$

As the above calculation illustrates, the Storage Service Unit portion accounts for about three times as much of the service as does the CPU Service Unit portion in the example.

The *Workload Manager: Planning* document notes that a large MSO value can cause SRM variables to overflow. In fact, the SRM monitors certain variables to determine whether they are in danger of overflowing. A long-running or never-ending "transaction" is stopped and restarted if the SRM determines that variables are in danger of overflowing. A large MSO service coefficient is the primary cause of variables overflowing.

- The Storage Service Unit values can be unpredictable in systems with significant central storage. This situation can arise because of varying demands for central storage.

  - Suppose that on one occasion, a particular address space executes in a workload mix in which central storage is heavily used. The central storage allocated to the address space would regularly be restricted to the basic working set of the address space. During this execution, the address space would have a service rate which might correspond to the original intent of the service rate computation.

  - However, suppose that on another occasion, the address space executes in a workload mix in which central storage is not heavily used. The central storage allocated to the address space might include many relatively "old" page frames - rather than simply the basic working set of the address space. During this execution, the address space would have a service rate that could be many times the service rate in the hypothetical situation in the preceding paragraph.

Under some circumstances, this situation would not be harmful. However, the SRM might place the address space into a lower service class period, with the performance goals and importance associated with the new service class period.

If the service rate is unpredictable and is unrelated to the actual requirements of the address space, then the address space will have unpredictable access to system resources. This unpredictable access is contrary to the basic intent of the management philosophy behind the design of the SRM and of the Workload Manager.

The following example illustrates the output from Rule WLM007:

```
RULE WLM007:  THE MSO SERVICE COEFFICIENT MAY BE TOO LARGE

   The MSO service coefficient was specified as 3, and CPExpert believes
   that this value is too large.  The MSO service coefficient is used to
   adjust the storage service portion of the overall service attributed by
   address spaces and to service classes.  The algorithm used by the SRM
   was designed when central storage was a major constraint in computer
   systems.  With increasingly large amounts of processor storage, central
   storage is not usually a major constraint to system performance.
   Consequently, installations often should reduce the  contribution that
   storage service makes to the overall computed service associated with
   address spaces and service classes.  Further, a large MSO coefficient
   used in modern systems can cause SMF or SRM fields to overflow, with
   resulting loss of data or erroneous actions by the Workload Manager or
   the SRM.
```

Suggestion: CPExpert suggests that you verify the specification for the MSO service coefficient in the service definition being analyzed.

If you are executing with a large central storage, you may wish to reduce the MSO service coefficient to zero. Reducing the MSO service coefficient will reduce the effect of central storage page frames on the service rate associated with an address space.

If you do change the MSO service coefficient, please review the DUR values you have specified for any multiple-period performance groups to make sure that the DUR values are appropriate with the decreased emphasis on storage service.

Reference: MVS Planning: Workload Management

MVS/ESA(SP 5): Chapter 12: Defining Service Units and Coefficients

OS/390 (V1R1):	Chapter 12: Defining Service Units and Coefficients
OS/390 (V1R2):	Chapter 12: Defining Service Units and Coefficients
OS/390 (V1R3):	Chapter 11: Defining Service Units and Coefficients
OS/390 (V2R4):	Chapter 11: Defining Service Units and Coefficients
OS/390 (V2R5):	Chapter 11: Defining Service Units and Coefficients
OS/390 (V2R6):	Chapter 11: Defining Service Units and Coefficients
OS/390 (V2R7):	Chapter 11: Defining Service Units and Coefficients
OS/390 (V2R8):	Chapter 11: Defining Service Units and Coefficients
OS/390 (V2R9):	Chapter 11: Defining Service Units and Coefficients
OS/390 (V2R10):	Chapter 11: Defining Service Units and Coefficients
z/OS (V1R1):	Chapter 11: Defining Service Units and Coefficients
z/OS (V1R2):	Chapter 11: Defining Service Units and Coefficients
z/OS (V1R3):	Chapter 11: Defining Service Units and Coefficients
z/OS (V1R4):	Chapter 11: Defining Service Units and Coefficients